SAMAN: Simulation Augmented by Measurement and Analysis for Networks

Ashish Goel

PIs: John Heidemann, Deborah Estrin, Ramesh Govindan, Goel
Students: Kun-chan Lan, Xuan Chen, Debojyoti Dutta
USC/ISI and UCLA

SAMAN Challenges

- *Network robustness* is a key challenge facing the Internet:
  - Understanding, predicting and avoiding failures
  - Hard to obtain accurate models, given the diversity of protocols and applications
  - Need to explore large search spaces to detect potential failures

Example Scenarios

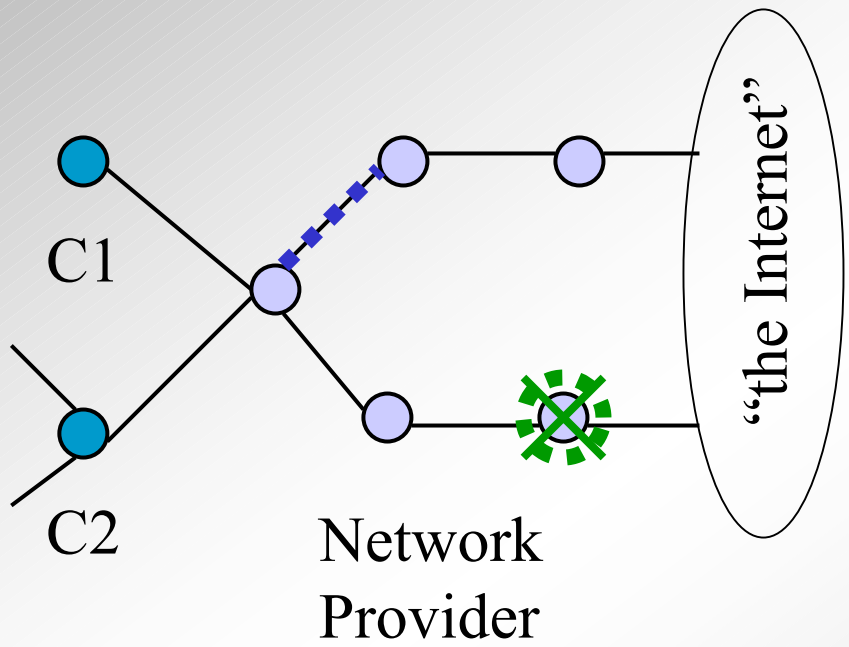

- The blue link becomes overloaded
  - SAMAN will help identify the cause

  => Need good traffic models
- What impact will the green router failure have?
  - "What if" scenarios

  => Need to explore correct part of large simulation space

Agenda

- Challenges
- **Analytic tools for pre-filtering simulations**
- Model generation – an example
- Understanding cascading phenomena
- Collaborations

Why Use Analytic Tools?

- Packet level simulators are very accurate but are time consuming
- Often, large chunks of the simulation space are uninteresting
  - Obviously bad, or obviously good
- A fast, approximate analytic pre-filtering tool could weed out uninteresting scenarios
  - Do detailed simulations for the interesting scenarios

Our solution at a glance

- Develop an analytic pre-filtering addition to NS (Approx-sim)
- Rapidly find approximate network operating conditions
  - Hybrid queuing theory approach with TCP equation [Padhye et al]
  - Order of magnitude faster than packet-level simulation
  - Approximate answers only (within 10% for symmetric trees)
- Use tool to select interesting scenarios for detailed simulation
- Currently prototyped for limited topologies and traffic (bulk TCP)

Overview of our architecture

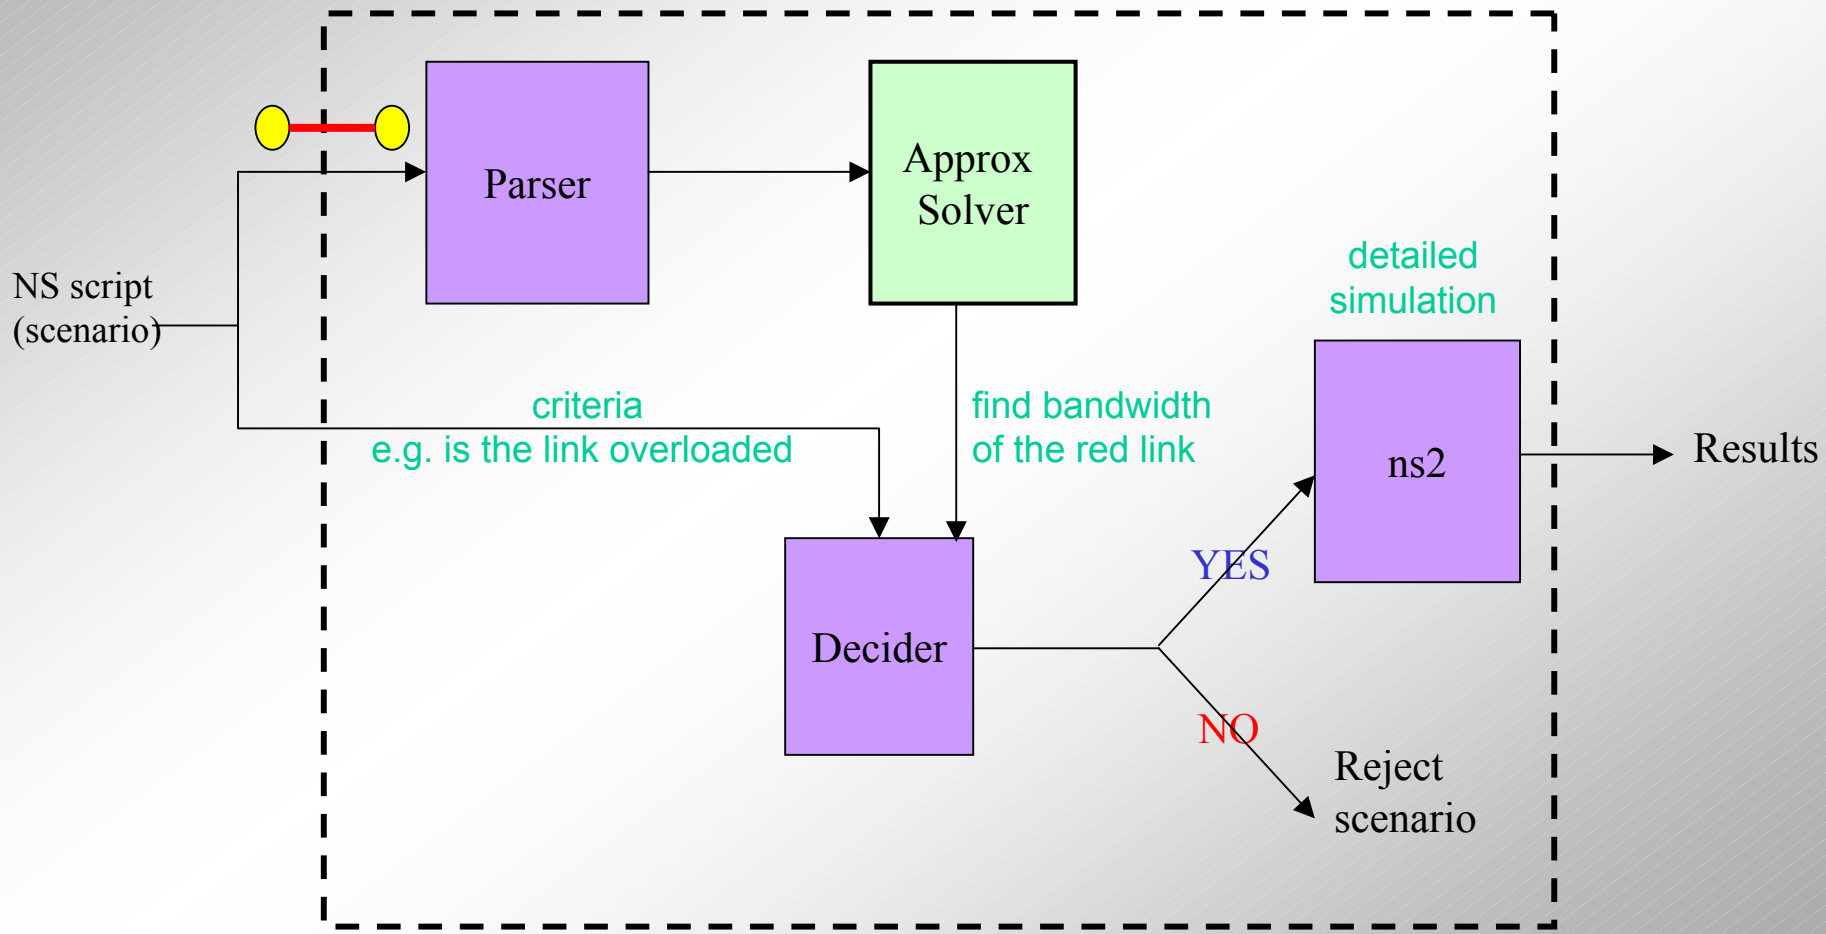

The basis for Approx-Sim

- TCP connections give rise to traffic on links
- Queueing theory gives approximate link characteristics (delay, drop probability), given TCP window sizes
- The TCP equation yields the approximate TCP window size, given link characteristics
- **Premise: The fixed point gives an approximation to the operating conditions**

The Fixed Point

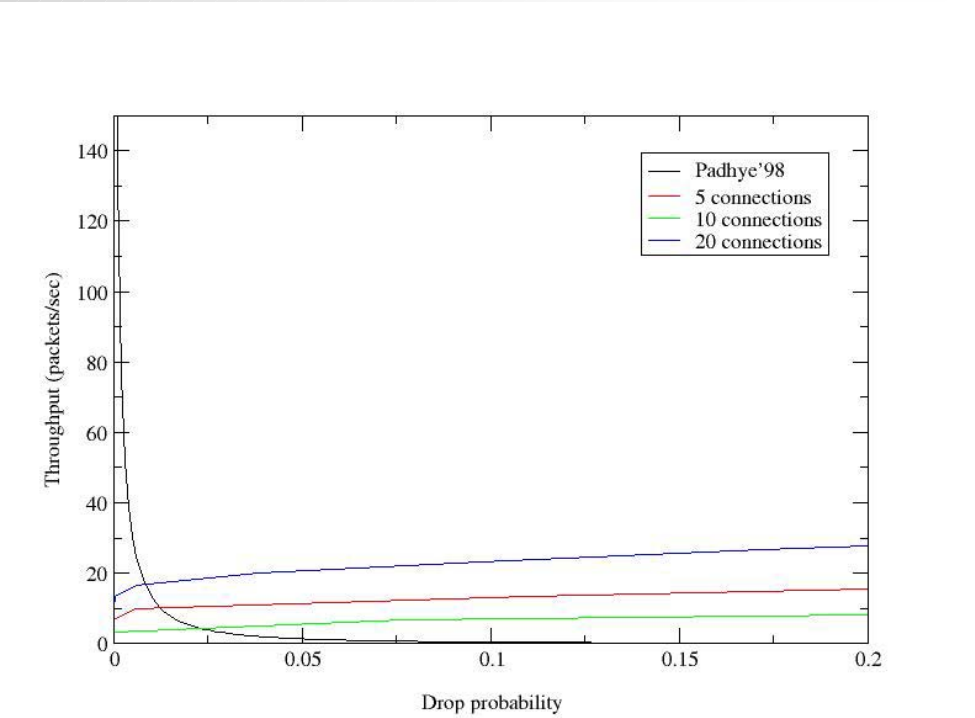

Theorem: There exists a unique fixed point for single router

Approx-sim under the lens

- Start with an estimate of throughput of TCP connections
- Calculate the drop probability and the delays using Queueing theory (M/M/1)
- Calculate end-to-end drops and delays
- Use TCP equation to recalculate the throughput

Naive Method Diverges

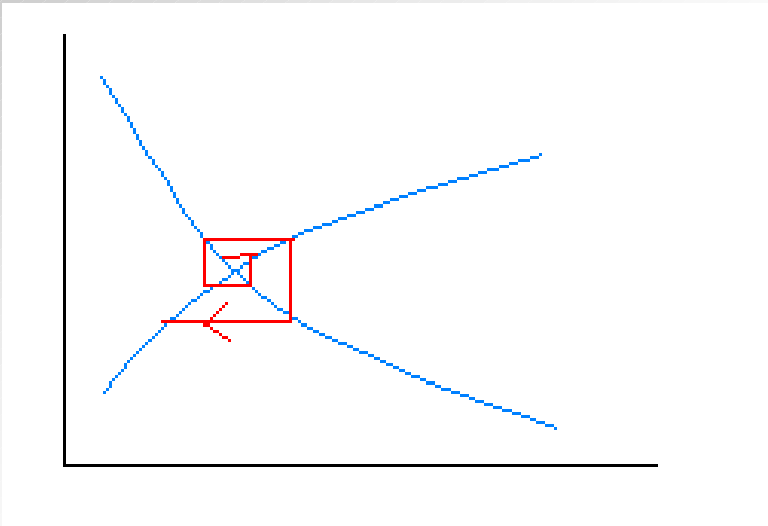

- TCP equation is concave, the queueing equation is convex
  - Simple iterative method does not converge
- Method would converge if we could run it in reverse
  - But equations not invertible

Main Problem: How do we ensure Convergence ?

Achieving Convergence

- For each overloaded link
  - Compute the overload factor
  - Scale throughput by this factor for each connection through this link
- Use new link throughputs to compute link characteristics (delay,drop probability)
- The TCP equation remains unchanged
- This technique has been tested for simple networks (lines, trees)

Approx-Sim is fast

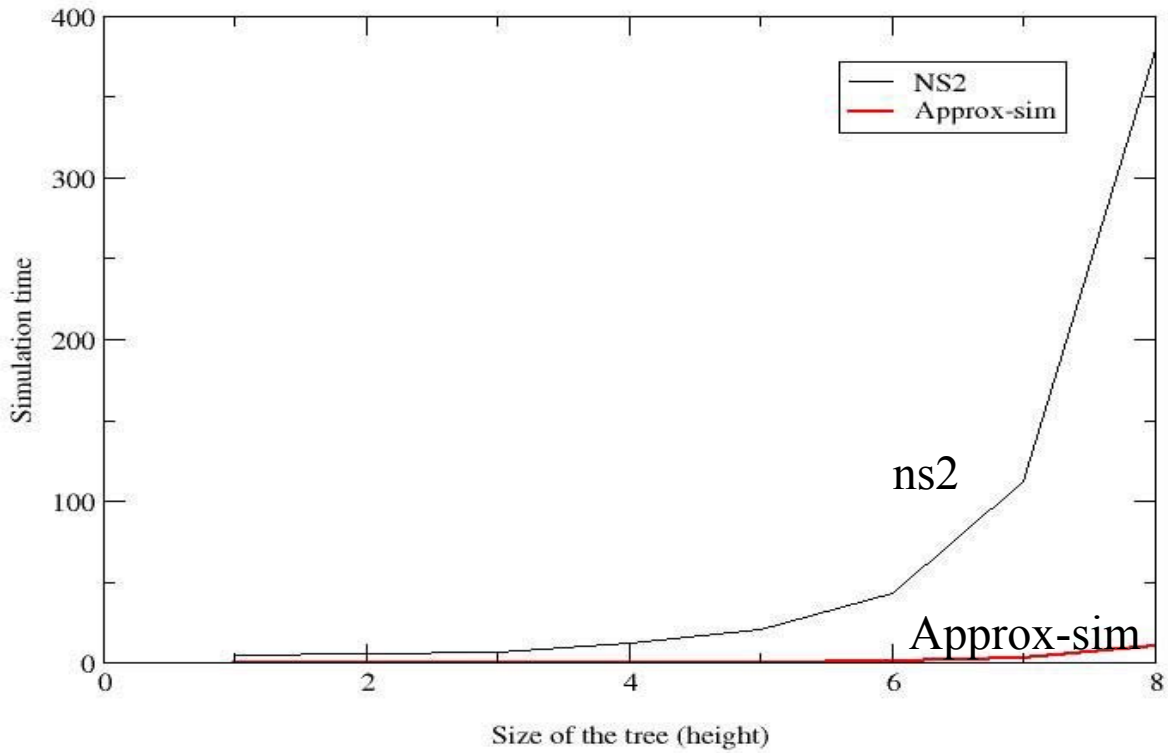

Accuracy counts

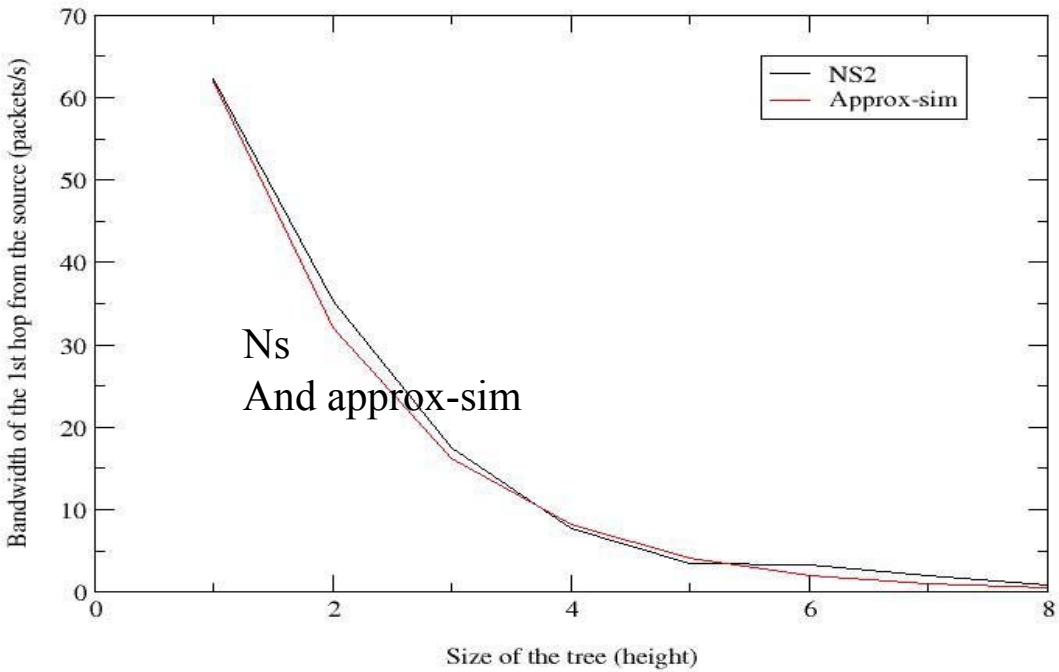

(Symmetric binary trees)

Future Work for analytic pre-filtering

- Generalize to support wider range of technologies (eg. RED) and test on general topologies
- Incorporate short TCP flows, UDP/CBR traffic
- Integrate Approx-Sim more tightly into NS
  - Have a simple query mechanism
  - Provide details of the simulation to user from the Tcl interface
- **Have done preliminary work on each of the above**
- Related Work: Bu, Towsley [Sigmetrics '01]

Agenda

- Challenges
- Analytic tools for pre-filtering simulations
- **Model generation – an example**
- Understanding cascading phenomena
- Collaborations

Multi-scale Analysis and Validation

- Demonstrate multi-scale analysis as a tool for validating and debugging traffic model
- Case Study : RealAudio
  - Develop structural model of RealAudio using traces
  - Validate using multi-scale wavelet analysis
  - Plausible explanation for the periodic burstiness of RealAudio
- Recent developments:
  - Revised validation
  - Protocol improvements inspired by model
  - Preliminary work at other traffic models

Model Validation using Wavelet Analysis

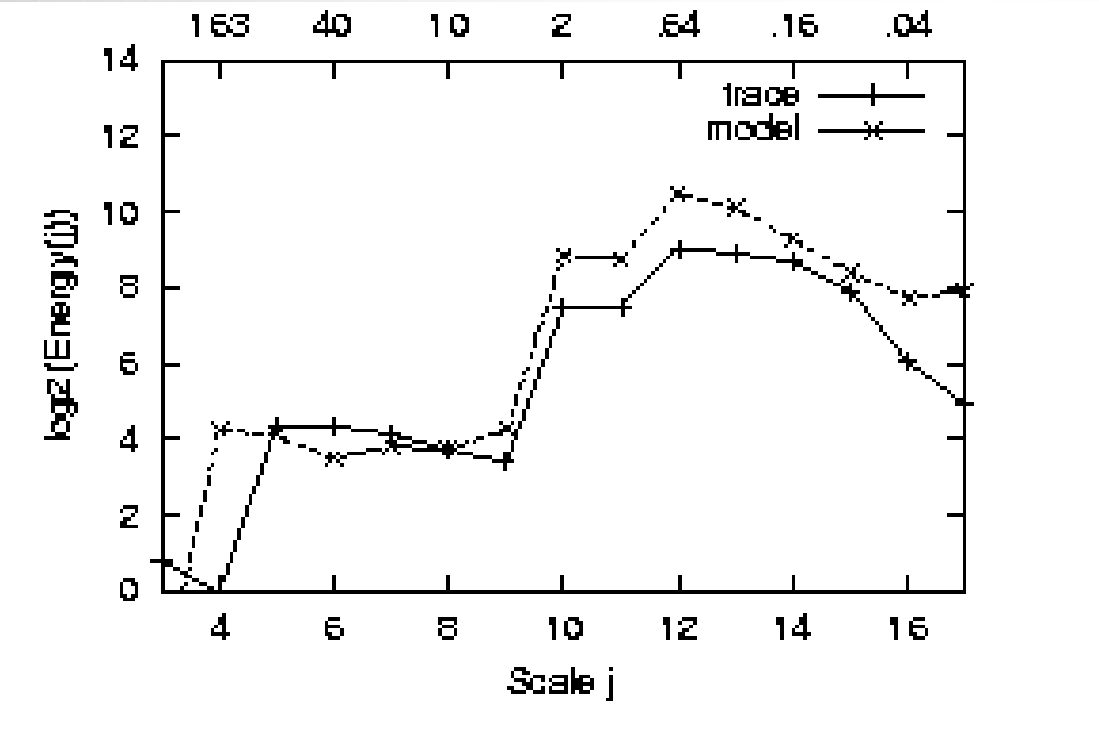

Performance Improvement after Minor Protocol Change

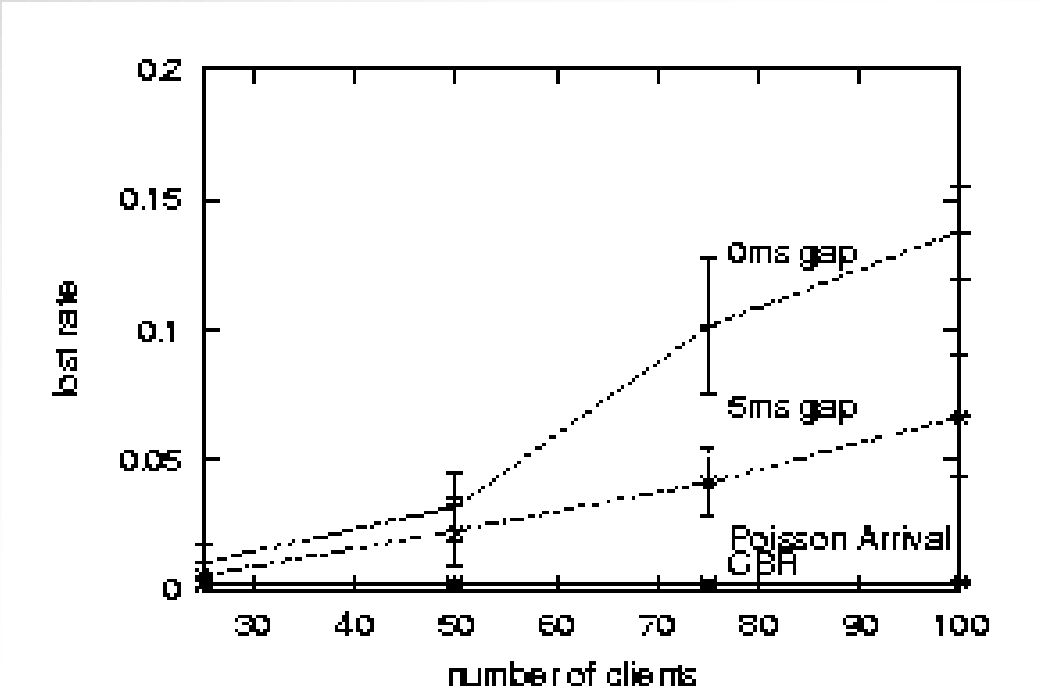

Minor protocol change to reduce burstiness

Ongoing Work: Real-time Model Instantiation

- Motivation: need to *quickly* parameterize models (minutes, not months)
- Issues
  - Model invariants independent of protocol (ex. User arrival patterns)
  - Integrate data measured at multiple points
  - How to automatically instantiate model from measurements
  - How to quantify model accuracy

Agenda

- Challenges
- Analytic tools for pre-filtering simulations
- Model generation – an example
- **Understanding cascading phenomena**
- Collaborations

Cascading Queues in Simulation

- Goal
  - Understand, detect, predict, and avoid network failures
- Current stage
  - Case study: cascading network phenomena
  - Understand cascading network failures from network traffic perspective

Methodology

- Reproduce cascading phenomena in ns simulation scenarios

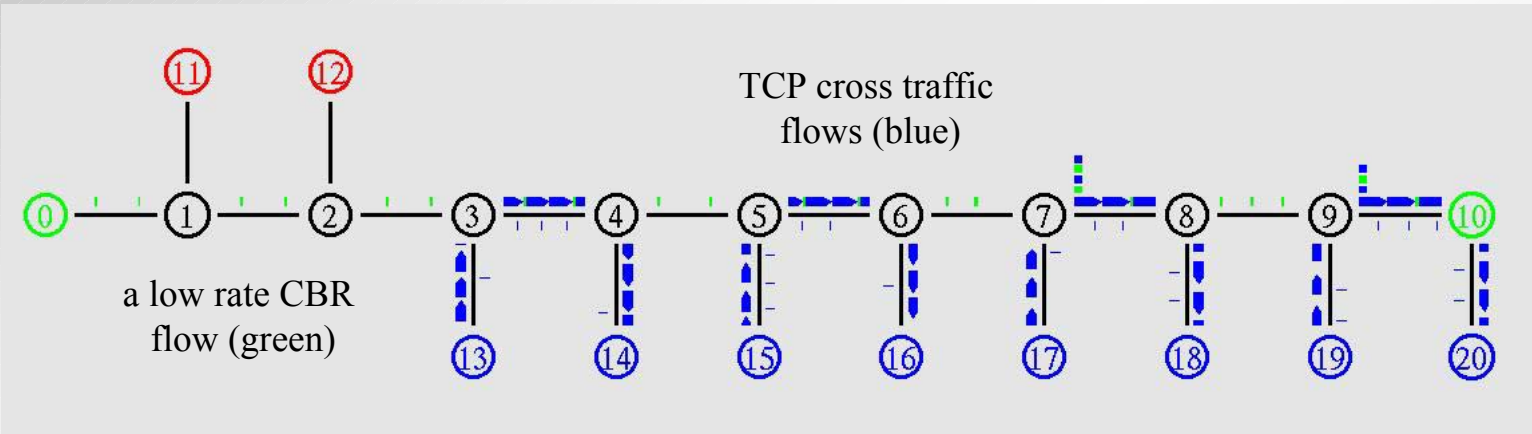

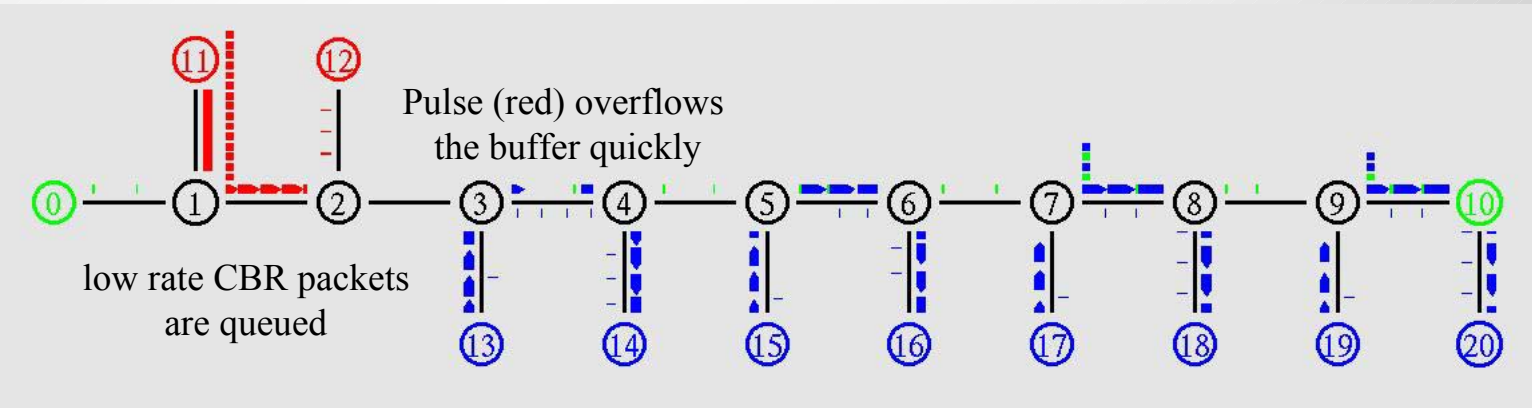

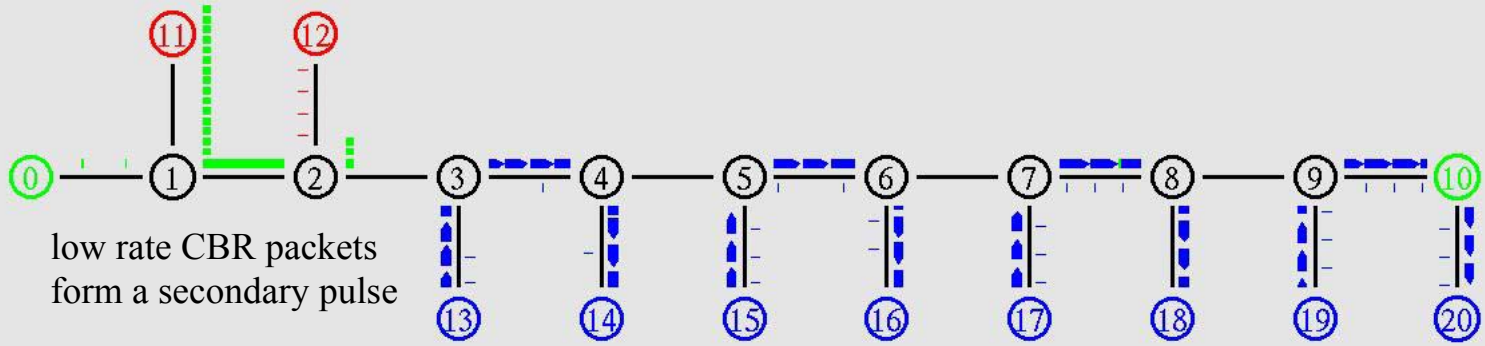

Secondary pulse (green) propagates and hits the TCP flows (blue)

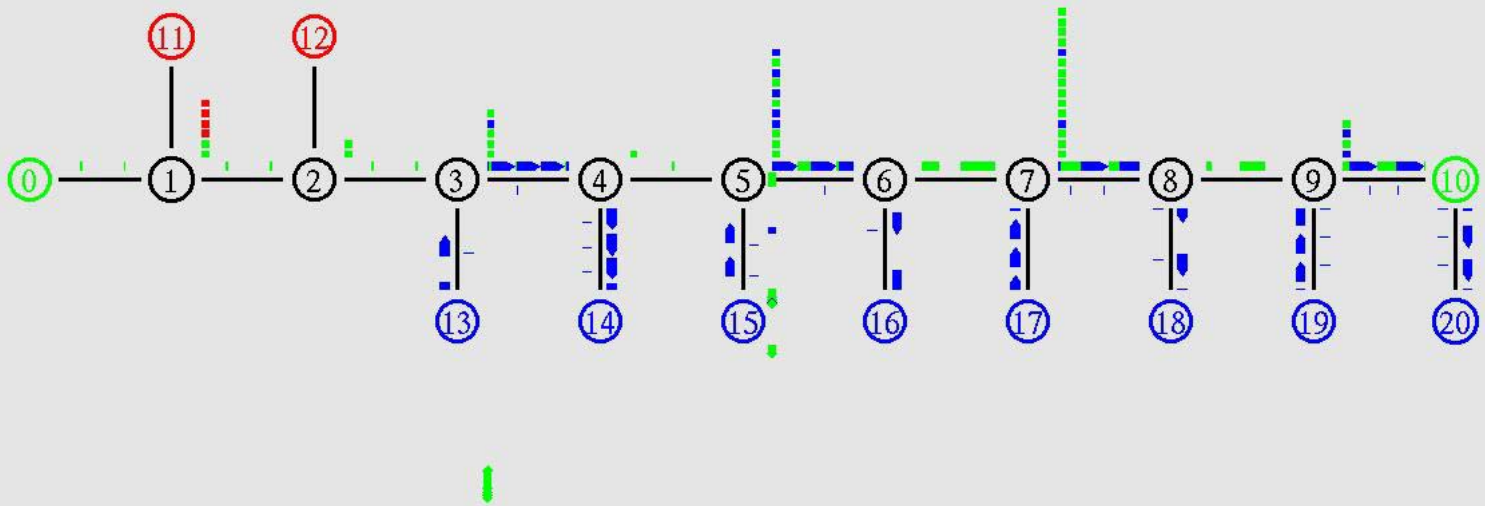

Status and Plans

- Have demonstrated cascading behavior in simulation
- Define indicators that characterize failures
  - packet loss rate
  - dynamics of the TCP round-trip time (RTT) and throughput
- Produce an early warning system to detect and predict network failures

Collaborations

- Cal Tech
  - Joint work to investigate cascading phenomena
- GaTech and RPI
  - Continuing integration of work into ns core
- CAIDA and other data providers
  - Use of network data to validate modeling